

In This Issue

CHARLES STELZRIED AND MICHAEL KLEIN

Data Services

In this issue's lead article, Sam Petty, et al., document the critical importance to the DSN of cryogenically cooled microwave low-noise amplifiers (LNAs). Improved LNA cryogenic cooling systems and the inclusion of critical feed components within them will lower operating noise. LNAs will continue to utilize both masers, for the very lowest noise performance with high RFI immunity, and high electron mobility transistors with excellent noise performance and wider bandwidth applications.

The new 1-m Optical Communications Telescope Laboratory (OCTL) is described by Nasser Golshan, et al. OCTL will be capable of tracking and communicating with laser-bearing spacecraft at distances from LEO to deep space during either day or night. OCTL will provide a unique R&D facility dedicated to support development of multi-gigabit communications from LEO satellites and tens to hundreds of kilobits-per-second downlinks from deep-space probes.

It is fairly well known that phase noise degrades the performance of deep-space communications systems. Jon Hampkins shows the advantage of coupling the DSN ground station receiver and decoder together as a partial countermeasure against phase noise. Currently, the DSN separates receiving and decoding functions. Improved performance is a result of the decoder refining the receiver carrier phase recovery and the receiver output refining the data decoding.

The goal of the Atmospheric Visibility Monitoring program is to quantify the effect of the atmosphere on optical communications. Babak Sanii describes this program, which consists of three autonomous observatories, each operating at three optical wavelengths. Atmospheric attenuation data is reported and compared with nearby surface weather observations. This information will be used for site selection and optical link margin calculations.

Mission Services

Tien-Hsin Chao describes a compact high-speed Grayscale Optical Correlator for real-time pattern recognition for spacecraft navigation applications. During a landing descent, when the target is matched with the correlation filter, a bright spot will be detected at the centroid of the target location. This system has demonstrated a 20:1 reduction of data processing compared to current methods. A smaller system is under development with anticipation that industrial partners will investigate packaging for a space environment.

Dan Lyons describes his early work on the effects of atmospheric drag on an object in orbit and its application to aerobraking for deep-space missions. Aerobraking is a low-cost and robust method for achieving orbit around a planet and has been used manually, but manual operations require a heavy dependence on continuous DSN coverage. Automating this process is under investigation for future missions.

Roberto Manduchi, et al., propose a solution for spacecraft instruments that collect more data than can be transmitted back to Earth. Visual techniques used to choose the most scientifically important data involve computing and then selecting from among multiple key elements. This technology addresses challenges in 3-D viewing of geologic features, development of fast algorithms, and incorporation of multiple views, all in order to create a robust on-board classification system.

Science

George Resch, et al., bring us up-to-date on the status of Water Vapor Radiometers (WVRs) in the DSN. WVRs remotely sense time-variable concentrations of water vapor in the atmosphere. The results are used to calibrate the propagation path delay of radio waves along the "line of sight" through the atmosphere toward a spacecraft or radio source. Possible applications include gravitational wave experiments, radio astronomy, and very long baseline interferometry. 🦊


DSN Low Noise Amplifiers In The New Millennium

J. BAUTISTA, R. CLAUSS, S. PETTY, J. SHELL

INTRODUCTION

The quality and volume of data collected from distant spacecraft can be enhanced by increasing the received signal strength or by reducing the noise in the Deep Space Network's (DSN) receiving systems. Cryogenically cooled microwave low-noise amplifiers (LNAs) have been used

since 1960 to reduce DSN receiving system operating noise to the lowest levels possible in an affordable and reliable way. Future improvements planned for this new millennium will provide LNAs with even lower noise, higher reliability,



CONTINUED ON NEXT PAGE

and more flexibility. Specifically, these improvements include lower physical operating temperatures, more efficient, reliable, and powerful cryogenic cooling systems, new amplifier designs and new active materials, critical feed-system components added within the cryogenic cooling system (to further reduce noise), and the support of different frequency bands in the same cryogenic LNA package.

Two types of cryogenically cooled LNAs are used in the DSN: masers and HEMTs (high electron mobility transistors). These LNAs are the first stage of amplification in the antenna's receiver chain, and are a crucial factor in obtaining ultra-low-noise receiving system performance. The HEMTs and masers compliment each other, each offering distinct advantages for various mission and science needs. The wide bandwidth and excellent stability of HEMTs enable the use of wide-bandwidth interferometric techniques for spacecraft navigation and provide the highest sensitivity for continuum radio astronomy. HEMTs also utilize reliable, commercially available two-stage cryogenic refrigerators operating in the temperature range of 5 to 15 K. Masers offer the lowest possible noise temperatures, a higher level of immunity to RFI, and a robust quality that defies any danger of burnout—an important consideration when 20- to 400-kW transmitters are on the same antenna.

This article describes the evolution of LNA performance during the previous millennium up to and including the latest LNA designs currently being developed or implemented into the DSN. The cost-savings rationale for developing and using these new LNAs is also discussed.

By convention, the amount of noise in an LNA is defined as its effective input noise temperature (T_e) and is measured in kelvins. The antenna noise temperature (T_a) includes noise from the cosmic background, the atmosphere, ground-pickup, and from reflector and feed system component losses. The total system operating noise temperature (T_{op}) is the sum of T_a and T_e , which includes noise contributions from the receiving system following the LNA. This notation (T_{op} and T_e) will be used in the following.

MASERS

The first cryogenically cooled LNA used in the DSN was a ruby maser installed just over 40 years ago, in September 1960. It used a laboratory-grown ruby crystal as the active amplifying material, operating in a 4.2-K liquid helium bath in

an open-cycle dewar, and demonstrated the technology on a DSN antenna at 960 MHz. A 2388-MHz maser was developed and used for the first successful Venus radar contact (1961). Single- and dual-cavity masers were developed for solar system radar experiments and for tracking the Ranger missions to the Moon. Traveling-wave masers (TWMs) were then developed to replace these early cavity masers; the limited bandwidth and tuning ranges of the cavity masers did not meet the growing needs of the DSN.

Many organizations participated in the development of masers between 1955 and 1965. The University of Michigan produced the first ruby maser. Workers at Bell Telephone Laboratories (BTL) and at the Airborne Instruments Laboratories (AIL) developed superb TWMs. TWMs were purchased for the DSN from AIL and their first use was to receive signals from the Mariner IV spacecraft to Mars during the Mars flyby in July 1965. These TWMs were cooled to a physical temperature of 4.5 K in closed-cycle refrigerators (CCRs), called Cryodynes by the manufacturer.

Industry soon realized that the use of masers (with their need for cryogenic cooling) was not cost-effective for communications from near-Earth satellites to Earth. It made good sense to use higher transmitter power on the Earth-orbiters and less-sensitive, low-cost receivers in the many ground-based terminals. By 1966, JPL could no longer depend on other organizations to develop the additional masers needed for future deep space telecommunications.


Dr. Walter H. Higa and Ervin R. Wiebe developed a 4.5-K closed cycle refrigerator with improved reliability and cooling capability in 1966. The mean-time-between-failure (MTBF) was increased from 1500 to 3000 hours and the reserve capacity at 4.5 K was increased from 700 mW to 1 W. Many of these CCRs are still in use in the DSN today.

Today, S-Band (~2.3 GHz) and X-Band (~8.4 GHz) masers in the DSN provide the lowest noise performance ever achieved for a TWM in a 4.5 K physical temperature CCR: four S-Band units, measuring 2 K noise temperature, and nine X-Band units, measuring about 3.5 K noise temperature, are in use (unless otherwise specified, LNA noise temperature is measured at the room temperature input connector on the cryogenic package). The X-Band TWM is shown in Figure 1. Three R&D masers are also in use. These systems

CONTINUED ON PAGE 24

AUTONOMOUS ATMOSPHERIC VISIBILITY MONITORING FOR OPTICAL COMMUNICATIONS

BABAK SANII



INTRODUCTION

Among the challenges of optical communication lies the significant hurdle imposed by nature: poor weather. Unlike radio wavelengths, which pass through cloud cover with little to no effect, a heavily overcast day is virtually opaque at the wavelengths of interest for optical communication. This is mitigated by the concept of site diversity; that is, ground stations may be strategically located in anticorrelated weather cells so that the statistical likelihood of transmission is increased by virtue of dissimilar weather. Nonopaque atmospheric attenuation is variable from day to day, site to site, and wavelength to wavelength. This is due to factors such as variations in thin Cirrus clouds, aerosols, molecular composition, density and temperature. The Atmospheric Visibility Monitoring (AVM) program, managed by the JPL Optical Communication Group, quantifies statistical atmospheric transmissions experimentally and evaluates the advantages of two- and three-site diversity. This information is critical to the determination of the locations and quantity of ground-based transceivers, as well as the link margin needed to establish successful optical communications through the atmosphere.

THE ATMOSPHERIC VISIBILITY MONITOR

The program consists of three autonomous observatories and a database of atmospheric transmission data

taken from the observatories. Each of the three observatories (fig. 1) is stationed in the southwestern United States; they are located at Mt. Lemmon (near Tucson, AZ), the Table Mountain Facility (near Wrightwood, CA), and the Goldstone Deep Space Communications Complex (near Barstow, CA). The three sites have been upgraded¹ and are now gathering data using new processors, software and cameras with greater near-infrared sensitivity and dynamic range. The sites are also networked to an atomic clock gateway, ensuring that they are taking images at the same time, permitting accurate site-diversity statistics to be tabulated.

CONTINUED ON NEXT PAGE

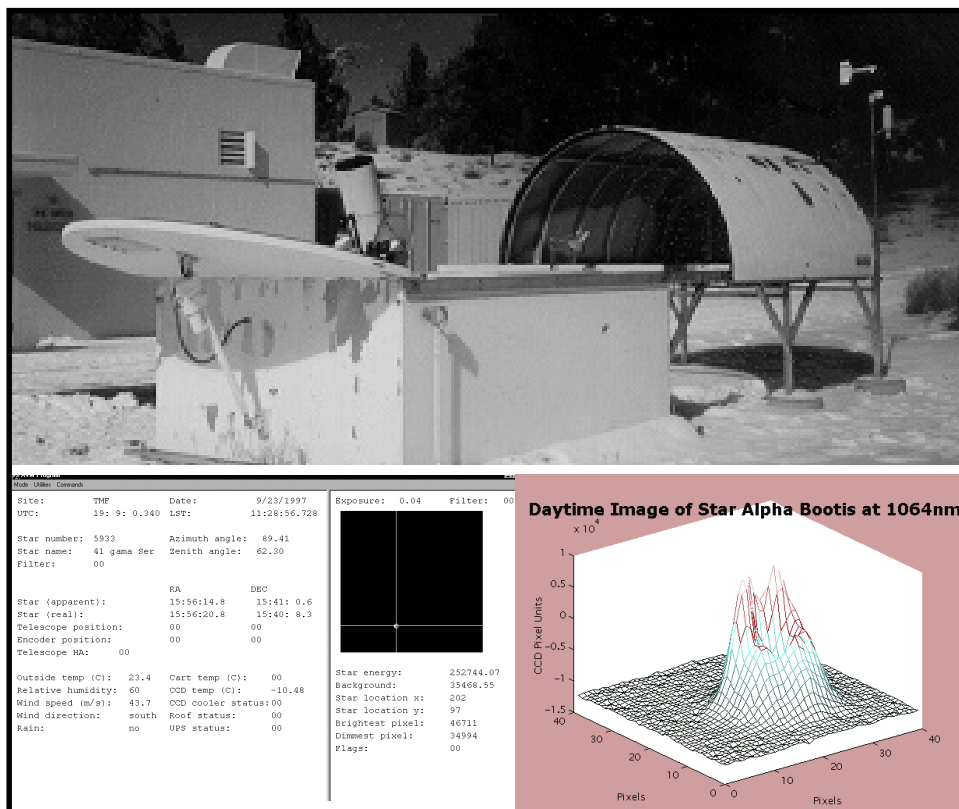


Figure 1. The AVM at Table Mountain (top); a screen shot of the software and a sample of a daytime image taken at 1064 nm (bottom).

Using motorized autonomous telescopes, the observatories take photometric readings of a set of bright stars through six optical filters. Three of these filters pass a narrow bandwidth of light that corresponds to laser wavelengths that may be used in future free-space optical communications systems, including 532 nm (a frequency doubled Nd:YAG laser), 1064 nm (the fundamental Nd:YAG wavelength), and 860 nm (a common diode laser). Measurements at the popular 1550 nm band are currently not supported by AVM. In addition to these, there are three other broadband astronomical filters in place: the standard Visual, Infrared and Blue filters. The observatory is completely automated, with environmental controls and alarms, and is capable of both day and night observations.

The program is an ongoing one. Since the upgrade, less than one year of data has been accumulated, so seasonal variations are not yet completely documented. Furthermore, several years of study will be needed in order to account for effects, such as El Niño and the solar cycle. In the meantime, some interesting preliminary results may be determined from the data.

OPTICAL WAVELENGTH ATTENUATION STATISTICS

Three laser wavelengths of interest are compared in the figure below for data gathered in May and June of 2000. In particular, by isolating data points where the atmospheric attenuation was less than 3 dB, histograms of zenith normalized attenuation are presented.

Figure 2 illustrates the statistical advantage of the 1064 nm wavelength. The 1064 nm line (squares) has a lower average attenuation than 860 nm (diamonds) or 532 nm (circles), as evidenced by the leftward horizontal shift of the histogram. The ordering of the statistical advantage is in accord with typical atmospheric transmission models such as MODTRAN², though the actual values are offset. This is particularly true at 532 nm, where the difference is almost 0.5 dB. The discrepancy may be due in part to a lack of any sort of record of detailed atmospheric composition, temperature, and altitude variations over time; these variables are needed to correctly implement MODTRAN to produce comparable results. While AVM is not intended to provide all of these variables, it records the principal result of interest for optical communication: transmission.




Figure 2. Histogram of AVM-reported attenuation at 532, 860 and 1064 nm

ASSIGNING ATTENUATION VALUES TO NATIONWIDE QUALITATIVE SURFACE WEATHER OBSERVATIONS

AVM data also has applications in describing other atmospheric observation tools with greater detail. For example, the National Climatic Data Center (NCDC) keeps hourly records of surface weather observations from hundreds of sites in the United States, particularly near airports. One of these surface observations is a rating of the current cloud cover, divided into four categories: Clear (no clouds visible), Scattered (1/8 to 4/8 of the sky is covered with clouds), Broken (5/8 to 7/8), and Overcast (complete cloud coverage). Furthermore, the Table Mountain facility is relatively near (about 40 nautical miles) to one of these observation sites (at Edwards Air Force Base) and is likely to be experiencing the same high cloud cover.

With this information in hand, the AVM data may be time-matched to the Edwards data, and attenuation distribution may be associated with different sky ratings. An example of this is contained in figure 3, using the average of the attenuation distributions for the first half of the year 2000. Further work validating this result is needed, but in essence this may allow a means of


assigning average attenuation statistics to any surface weather observation site reporting to the NCDC.

Work is currently underway performing a similar analysis using satellite imagery of the AVM sites. This may enable worldwide analysis of atmospheric visibility and ground station site-selection.

ACKNOWLEDGMENTS

The work described was funded by the TMOD Data Services Technology Program. Ana Datta, Beckett Madden-Woods, and Abhijit Biswas contributed to this article.

REFERENCES

- [1] B. Sani, A. Datta, D. Tsiang, J. Wu and A. Biswas "Preliminary Results of an Upgraded Atmospheric Visibility Monitoring Station" *TMO Progress Report 42-142* 2000.
- [2] A. Berk, L. Bernstein, G. Anderson, P. Acharya, D. Robertson, J. Chetwynd, and S. Adler-Golden "MODTRAN cloud and multiple scattering upgrades with applications to AVIRIS" *Remote Sensing of Environment* 1998. 





Figure 3. Bar graph of the average AVM reported attenuation at Table Mountain, corresponding to sky cover observations at Edwards. Qualitative sky cover observations are taken at locations nationwide by trained personnel



NASA's ONE-METER OPTICAL COMMUNICATIONS TELESCOPE LABORATORY

NASSER GOLSHAN, Ph.D., KEITH WILSON, Ph.D., MICHAEL BRITCLIFFE

INTRODUCTION

Technology developments, studies, and successful laser communications demonstrations toward the end of the last century have enabled laser communications researchers to gain a better grasp on the issues pertaining to reliable high-bandwidth space-to-ground optical links. This work has led to the development of plans for an aggressive suite of technology demonstrations in the 2001–2010 time frame. In response to the expected increase in the number and complexity of technology demonstrations, NASA/JPL is currently building an Optical Communications Telescope Laboratory (OCTL) at its Table Mountain Facility (TMF) near Wrightwood, California. This facility is designed to be a state-of-the-art optical communications laboratory capable of tracking and communicating with laser-bearing spacecraft at distances from low-Earth-orbit (LEO) to deep space. The facility will be outfitted with the instrumentation needed to support multi-gigabit communications from LEO satellites and tens to hundreds of kilobits per second downlinks from deep-space probes. Key activities being planned for OCTL in the years ahead include:

- Develop and demonstrate precision acquisition and tracking of objects from LEO to deep space
- Explore uplink beam propagation strategies, and define approaches to mitigate the effects of atmospheric turbulence
- Validate predictions of daytime optical reception and develop strategies to suppress the effects of sky background on the signal-to-noise ratio (SNR) of the deep-space downlink
- Explore strategies to mitigate atmospheric effects on space-to-ground coherent optical communications
- Design, assemble, and test optical transmitter and receiver concepts needed for future deep-space communications

The optical telescope system is designed to operate at any polarization in the 500–2200 nm wavelength range to accommodate typical operating frequencies for a number of likely optical communications experiments as shown in Figure 1. Among the first demonstrations planned for the OCTL will be to acquire and track retro-reflector bearing satellites such as are used for satellite laser ranging. Spacecraft of opportunity with optical retro-reflectors include Lageos 1 and 2, Etalons 1 and 2, TOPEX/Poseidon, and Ajisai. These will make excellent targets for validating the acquisition and tracking performance of the telescope. Early experiments will include a comparison of fade statistics of multiple- vs. single-beam propagation using candidate satellites that range from 1500–20,000 km altitude. Current CRL/NASDA (Communications Research Laboratory/National Space Development Agency of Japan) plans call for launching the OICETS (Optical Inter-orbit Communications Engineering Test Satellite) in the 2001–2002 time frame. Because this is a LEO satellite it will be visible from ground stations. We plan to explore performing an optical communications demonstration with this satellite.

Figure 2 shows a schematic of the telescope system. The design is based on a 7-mirror beam-waveguide system. The motorized M-7 mirror can be positioned in 90-degree increments to allow experiments to be conducted from four locations. This design is similar to RF beam-waveguide antennas but different terminology is used in the trade: the beam-waveguide design is called Coude Focus and the experiment room is known as the Coude room. Provisions have been made for the addition of a Cassegrain focus as a future option.

OCTL has been designed for both night and day operation. While performance specifications are guaranteed for Sun separation angles greater than 30 degrees, the telescope system is designed

CONTINUED ON PAGE 8

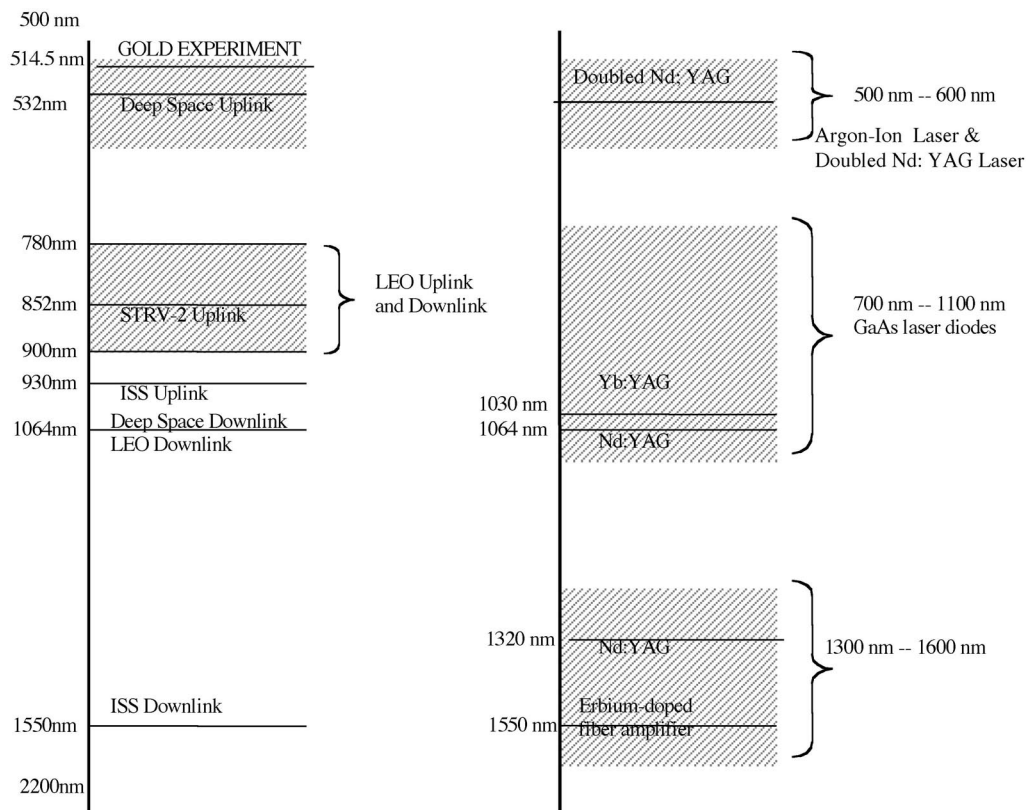


Figure 1. Typical operating wavelengths for a number of likely optical communications experiments; OCTL shown on the left, typical lasers and corresponding wavelengths are shown on the right




Figure 2. Schematic of the 1-m telescope with acquisition telescope and dome (shows pier, isolation pad, Coude path mirrors, and a laser located on an optical table near the Coude focus)

to survive when operating as close as 10 degrees from the Sun. Sun avoidance routines are incorporated into the control software to keep the telescope at least 10 degrees from the Sun. The total daytime scattered light from coated mirrors and window surfaces will be one order of magnitude lower than scattered sunlight from the portion of the atmosphere directly in the field-of-view (FOV) of the telescope.

A specially designed 200-square-meter facility has already been constructed at JPL's Table Mountain Facility to house the telescope (fig. 3). A contract has been competitively awarded to Brashear LP for the delivery and installation of the 1-m optical telescope system. The turnkey telescope system also includes a 20-cm acquisition aid telescope; a telescope tracking and pointing system; a dome; and associated hardware and software for control/operations/safety. Three major milestones, including Optics Design Review, Preliminary System Design Review, and Critical Design Review, have been completed; delivery and installation of the turnkey telescope system is planned for September and October of 2001, respectively. The telescope design is based on a 1-m telescope built by Brashear for the

United States Air Force (fig. 4). The design has been upgraded to attain high optical transmission efficiency and to limit stray daylight getting into the Coude focus, thus enabling detection of very weak optical signals.

PERFORMANCE SPECIFICATIONS

The telescope system is being designed for high energy-collection efficiency for optimal reception of weak signals. At least 60 % of energy incident on aperture will be received at Coude focus for $\lambda > 600$ nm, and greater than 50% for $600 \text{ nm} > \lambda > 500$ nm.

The need to form sharp images at Coude focus for efficient signal detection and tracking derives the wave front error (WFE) requirements of OCTL. When the optical telescope is pointed to the zenith and focused at infinity, the nominal WFE at Coude focus will not exceed 0.054λ , RMS, at 633 nm over the temperature range $+10^\circ$ to 20° C. This will ensure that at least 65% (85%) of the energy will be encircled in a $1.54 \mu\text{rad}$ ($2.81 \mu\text{rad}$) spot. The telescope system has an effective focal length of 75.8 meters, providing a $500\text{-}\mu\text{rad}$ flat FOV on a 37.9-mm diameter circle on a CCD imaging array.

Figure 3. OCTL building at Table Mountain facility.
Top, interior of OCTL before roof/interior partition installation, showing Coude isolation pad's center bolt pattern.
Bottom, building exterior after roof installation.






Figure 4. A 1-m telescope built by Brashear for the U.S. Air Force; the JPL telescope is based on this design but incorporates a number of upgrades

The optical telescope system has an integrated control system with a user friendly interface. This feature will allow a single operator to configure, control, and operate the telescope, dome, and pedestal mirror and drive either from the OCTL control room or remotely over the network. Configuration, status, and performance displays are being provided for all modes of operation. These displays permit a single operator to determine the operational and health status of the optical telescope system. Monitor data is provided on the tracking performance of the telescope. Such monitor data includes predicted pointing state-vector, commanded pointing state-vector, and servo tracking error. The telescope control system allows for automated logging of monitor data into a monitor file. The telescope control system allows the telescope and dome controls to be reset via either a local command or authorized remote commands. The telescope system can also be operated via a portable computer from inside the dome for maintenance or troubleshooting purposes. An emergency stop inside the telescope dome helps prevent accidental movement of the telescope when people are working on it.

The telescope pointing, tracking system, and its control system provide the capability to track satellites from LEO to deep space, and will also be able to point to fixed AZ and EL positions, such as would be used in a horizontal path demonstration or when pointing at a geostationary satellite. The pointing and tracking modes are: 1) Local mode where the telescope points to fixed AZ and EL positions (-5° to 185°); 2) planetary mode where the telescope tracks planets, asteroids, and other solar system bodies from a vendor-supplied catalog; 3) orbital mode where the telescope tracks objects in Earth orbit based on a NORAD two-line orbital element set; 4) sidereal mode where the telescope will develop a star file for mount calibration; and 5) predict mode where the telescope tracks objects from LEO to deep space from externally provided files containing state vectors (time-stamped AZ/EL values) and start/stop times. The telescope control system allows manual adjustment of telescope pointing during track to allow the satellite to be centered in the main telescope's FOV. The telescope control

CONTINUED ON NEXT PAGE

system will be able to update the satellite's ephemeris file for the pointing adjustment.

The telescope system will be able to image objects from 1 km to infinity at Coude focus and be compatible with laser beam uplinks. The Coude mirror coatings are optimized to support uplink peak power densities $> 10 \text{ MW/cm}^2$ in 10 ns pulses at 250 Hz, at 1064 nm, and at 532 nm, and average power outputs of 100 W. Special attention is paid to pointing accuracy to facilitate uplink experiments. The Coude room, the telescope, and its foundation have been designed to maintain blind pointing accuracy to within 15μ radians for both daytime and nighttime operations over the operating temperature range of -10°C to 40°C . This is achieved by incorporating an error correction all-sky model into the telescope control code for calibrating the telescope mount. For telescope travel within five degrees of a calibration reference, the line-of-sight will retain a relative accuracy of $2.5 \mu\text{rad}$. Telescope focus change with temperature will be calibrated, and the calibration results will be accessible to the operator. The M-7 mirror is motorized, remotely controlled, and allows 360° rotation; positioning is repeatable to within $10 \mu\text{rad}$.

These capabilities allow for flexible implementation of various acquisition and tracking strategies. For example, if accurate enough ephemeris information (GPS, NORAD, TDRS, or DSN predicts) is available, the telescope can use the ephemeris file to blind point, track, and uplink the beacon. If such information is not accurate enough for blind pointing but is sufficient to get the target image within the acquisition camera's FOV, the LEO target, if visible, will be imaged onto the acquisition camera and used to guide the main telescope. When errors in position knowledge are projected to be worse than acquisition camera's FOV, telescope predicts are loaded but telescope is put in a wait state until the satellite enters the FOV. At that time the satellite is centered in main telescope and then the ephemeris epoch is adjusted for time delay.

The need to track LEO spacecraft, the Space Shuttle (STS) in particular, drive the requirements on tracking rate, acceleration, and jitter requirements of the optical telescope system. The telescope system will be able to track at rates and


accelerations as needed for a nontracking keyhole no greater than a 6-degree radius for spacecraft in LEO (250 km altitude). The nontracking keyhole will be halved for each doubling of the spacecraft orbit altitude. The optical telescope system has a slew rate of at least 20 degrees/s in azimuth and 5 degrees/s in elevation. Telescope elevation travel limit allows pointing the telescope from -5° to 185° and to track from 0° to 180° , excluding the keyhole for LEO spacecraft. Telescope azimuth travel limit will allow pointing the telescope and tracking over $\pm 200^\circ$.

The telescope and all mirrors are designed to allow both quick removal of mirrors for recoating and easy reinstallation of mirrors without degrading the performance of the telescope system. In addition, the mirrors are being provided with handling fixtures and easy to follow instructions for removal, transportation, and reinstallation. These fixtures are compatible with a mirror recoating facility located in the contiguous United States authorized by the telescope supplier for proper recoating of the mirrors. The telescope and dome design allows easy cleaning of the mirrors in place.

CONCLUSIONS

A state-of-the-art 1-m Optical Communications Telescope Laboratory is being constructed by JPL for NASA at JPL's Table Mountain Facility near Wrightwood, California. The telescope vendor, Brasher LP, has completed all major design milestones; delivery and installation of the telescope system is planned for September and October of 2001, respectively. When completed, OCTL will provide a unique R&D facility for NASA dedicated to support multi-gigabit communications from LEO satellites and tens to hundreds of kilobits per second downlinks from deep-space probes.

ACKNOWLEDGMENTS

Research reported by this work is being performed by Jet Propulsion Laboratory, California Institute of Technology under a contract with National Aeronautics and Space Administration in support of optical communications technology for cost effective telecommunications support of NASA missions. 

COUPLED RECEIVER-DECODERS

JON HAMKINS

INTRODUCTION

A spacecraft communicates with Earth by propagating electromagnetic energy, much in the same way that a radio station transmits voice or music. Instead of analog voice or music signals, however, usually binary digits—zeroes and ones—are transmitted from a spacecraft.

The digital transmission can be accomplished by sending one electromagnetic wave to represent a zero and another wave with a 180-degree phase shift to represent a one. An extra electromagnetic continuous wave (a “residual carrier”) is also usually sent. The wave is just a sinusoid and contains no actual data, but it helps the receiver on ground synchronize the phase of the incoming signal so that the receiver can properly recognize the signal’s 0- and 180-degree phase shifts.

Signal power is balanced between the residual carrier and the data signal. Too much power in the residual carrier would mean that the carrier phase synchronization would be easily accomplished, but at the expense of being able to detect the 0- and 180-degree phase shifts in the lower power data signal. On the other hand, too little power in the residual carrier would mean that the carrier phase might not be able to be synchronized properly, and thus the frame of reference for measuring the phase shifts would not perform well enough to detect the binary digits reliably.

HOW DO WE DETERMINE WHAT DATA WAS SENT?


On the ground, the task of determining the binary digits transmitted is conventionally divided into two parts. The first task is to determine the

phase of the incoming signal. This is accomplished by the receiver, which uses a feedback loop called a phase-locked loop to process the residual carrier and produce an accurate estimate of the phase. Based on this estimate, the data signal is correlated with hypothetical “zero sent” and “one sent” signals; and the correlation coefficients, called soft symbols or matched filter outputs, represent the receiver’s output.

The second task is to determine the binary digits that best explain the matched filter outputs coming from the receiver. For a transmission of uncoded binary digits, this can be accomplished on a symbol-by-symbol basis by considering the sign of each matched filter output. In most systems, however, the data is coded. This means that redundancy is added to the binary digits before transmission so that if an error occurs in transmission it can be detected and corrected on the ground without the need for retransmission. On coded data, the decoder does not operate on a symbol-by-symbol basis. Instead, it considers a block or sequence of matched filter outputs in deciding the value of each output bit.

Traditionally, the tasks of carrier recovery and data decoding have been accomplished in the isolated manner described above. This architecture is shown in Figure 1. The sole communication between the receiver and decoder is via a sequence of matched filter outputs at the input of the decoder.

There is a good rationale for separating the receiver and decoder operations. Each of the phase and data recoveries may be optimal, individually.



CONTINUED ON NEXT PAGE

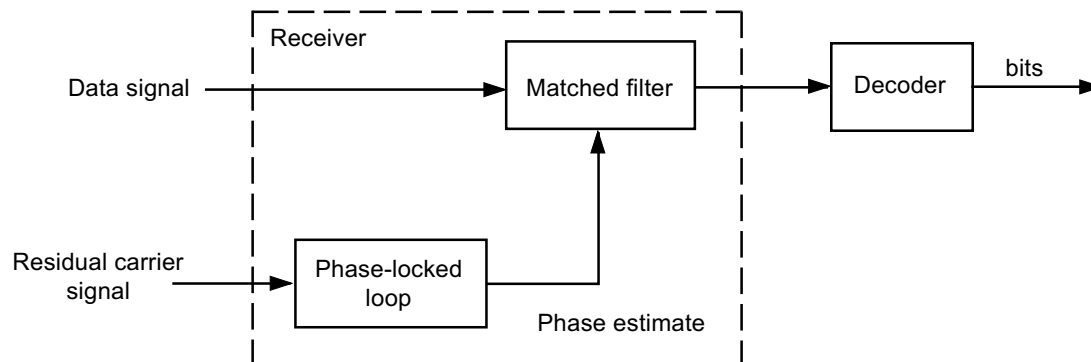


Figure 1. An uncoupled receiver and decoder

In fact, this may be proven mathematically in some cases. However, the optimality of phase tracking loops is based on the assumption that transmitted symbols are independent, and the optimality of a decoder is based on the assumption of perfect carrier phase recovery. Neither assumption holds for coded, phase noisy signals of the type usually seen in deep space communications.

What we have found works much better is a coupled receiver and decoder, where information transfer is bidirectional: The decoder output helps refine the carrier phase recovery, and the receiver output helps refine data decoding.

A JOINT RECEIVER-DECODER

The ultimate limit of coupling the receiver and




Figure 2. A joint receiver-decoder

decoder occurs when the phase and decoded data are jointly estimated. In this scheme, shown in Figure 2, the previously separate receiver and decoder operations are rolled into one. Hypothesized bit sequences are compared in the decoder using matched filter outputs to see which sequence was most likely transmitted, just as would happen in any decoder. The key innovation is that each bit sequence has its own associated phase estimate. This is crucial, because the best phase estimate under the hypothesis that one bit sequence was transmitted differs from the best phase estimate under the hypothesis that some other sequence was transmitted. The result is that the matched filter outputs for the correct bit sequence are of a higher quality than they would be in a separated architecture.

This joint receiver-decoder is at its best when used with what is called convolutionally encoded data. This type of coding

provides an efficient method of comparing sequences of transmitted bits that is amenable to the joint receiver-decoder architecture.

This architecture was tested on a NASA standard (7,1/2) convolutional code and a simulated DS-1 Ka-band link. Figure 3 shows the Bit Error Rate (BER) performance versus the total transmitted energy-to-noise ratio per bit. This is an apples-to-apples comparison across receiving architectures. In the uncoupled system, signal power is balanced between the data signal and the residual carrier signal in the proportion that results in the best BER performance. In the joint system, signal power is concentrated only on the data signal, a so-called suppressed carrier signal. Seemingly by magic, we observe a dramatic improvement in performance over the uncoupled system, even though both systems use the same total transmission power. The joint receiver-decoder system saves 3 dB at a BER of 0.01. If a genie were able to predict the exact phase of the incoming signal, another 2 dB could be saved—this is a bound on the performance of any system.

A COUPLED RECEIVER-DECODER

For some types of codes, there isn't a convenient maximum-likelihood sequence estimator as there is for convolutional codes. Turbo codes, for example, show enormous advantages over many other types of codes, but they use an iterative decoding algorithm that does not lend itself easily




Figure 3. Comparison of joint and uncoupled receiver-decoders for convolutional codes




Figure 4. A coupled receiver-decoder for turbo codes

to the joint receiver-decoder architecture. We can use the iterative nature of the decoder to advantage by explicitly passing information from the decoder to a separate, but coupled, receiver. This architecture is shown in Figure 4.

In this scheme, in the first iteration the receiver operates as it does in the uncoupled architecture, passing matched filter outputs to the decoder in the usual way. The turbo decoder goes through its first decoding iteration, in which it makes a rough cut at estimating the transmitted data and its reliability. The decoder then sends soft data symbols back to the receiver, where they are used to refine the phase estimates used in the first iteration. Thus, refined matched filter outputs are used in the second turbo decoder iteration. The process is repeated for all subsequent iterations.

Figure 5 shows the BER performance of the architecture for a rate 1/3 CCSDS standard turbo code, again on a simulated DS-1 Ka-band link. The leftmost curve illustrates the nominal performance of the turbo code when perfect carrier

reference is assumed, based on a suppressed carrier signal. At the other extreme, the rightmost curve illustrates the performance of an uncoupled system when phase noise is present. The proportion of signal power allocated to the residual carrier was chosen to optimize BER performance while meeting current receiver requirements. In-between is the curve for the coupled receiver-decoder. The residual carrier power was chosen to approximately optimize the BER performance. The carrier tracking loss is about 2.0 dB, which represents a 3.75 dB improvement over the uncoupled architecture.

CONCLUSIONS

This article discussed methods to couple receivers and decoders for phase noisy signals. The joint and coupled receiver-decoders were simulated on low-data-rate signals, where phase noise becomes most problematic, and were found to save 2.8 to 3.8 dB in total transmit power over simple uncoupled schemes. 📡





Figure 5. Comparison of coupled and uncoupled receiver-decoders for turbo codes



ONBOARD SCIENCE PROCESSING AND 3-D MODELING FOR SMARTER COMMUNICATIONS

ROBERTO MANDUCHI, BECKY CASTANO, AND ISAAC COHEN (USC)

INTRODUCTION

More data can be acquired with spacecraft and rover instruments than can be transmitted back to Earth. For example, the Athena Rover uses up to 12 cameras for trajectory planning and accurate surface mapping. The huge volume of data generated cannot be either stored or downlinked. This is the case for imagery acquired by the “hazcams,” which are cameras used solely for the purpose of obstacle detection. However, this data may still be used to extract features of geological interest, which can be encoded very efficiently and transmitted for the benefit of both the scientific community and the general public.

In an article that appeared in the previous issue of TMOD Technology and Science Program News [Keilly 2000], we presented a prioritized buffer management scheme, which ensures that the highest priority data segments gathered in a mission are always transmitted first, and only the least valuable data segments are ever discarded. Here, we will discuss how we can use visual techniques to analyze geological formations and then carefully select the most scientifically interesting data to return.

FIELD GEOLOGY MEETS DEEP SPACE COMMUNICATIONS

Progressive compression techniques have been developed for Earthbound applications with many of the same objectives that we seek to accomplish using deep space communications. However, deep space communication differs from Earthbound applications in one important respect. On Earth there is virtually unlimited archival storage, so valuable data never needs to be thrown away. In contrast, finite spaceborne storage resources will always force scientists to decide what data to keep and what data to throw away (or never collect). In order to *prioritize* the acquired data according to its scientific importance, we

need to make sensible use of the on-board information processing capabilities. In particular, we are working to develop automated field geology algorithms based on visual texture and three-dimensional surface representation.

What are some “interesting” geologic characteristics that could be used to prioritize the data acquired by a moving rover? For instance, we may seek features that indicate the former presence of water and, thus, locations with higher than average opportunities for finding signs of life—if any exists. Examples of such indicators are layered sedimentary deposits (identified by their visual appearance) and carbonate rocks (detected via their spectral characteristics.) Another criterion for determining scientifically interesting target information is the novelty of the data. It is desirable to sample each type of rock or formation encountered during a rover traversal. Thus, when a new rock or structure is encountered it should be given a high priority for transmission, while a rock type that has already been sampled and transmitted should be given a lower priority.

NASA has demonstrated a growing interest in autonomous field geology. Machine learning techniques for the analysis of visual texture features and spectral signatures have been developed and tested using data collected by the rover FIDO during field tests conducted in the Mojave desert [Gillmore 2000]. These techniques have shown promising results in terms of rock classification. Figure 1 shows a “rock finder” that separates rocks from soil based on simple color processing techniques. While most of the rocks have been correctly separated from the soil, some instances of missed detection remain due to the inherent ambiguity of color information. Other visual cues, such as visual texture and range information from stereo cameras, are currently being used to increase the robustness of the system [Gillmore 2000].



Figure 1. A simple color-based rock finder.
Dark areas represent detected rocks

ENTER 3-D

While single-view data prioritization allows us to increase the science return via “semantic” prioritized transmission, a great deal of redundancy is unavoidable when the same surface is imaged multiple times by a moving rover. There are two aspects related to such redundancy. First, if the same rock is imaged from two or more viewpoints, we may expect that its geological classification will be more robust and accurate than in the single viewpoint case. However, this means that we need a geometric scene description in order to correlate the different views with each other, as well as to correct the visual textural properties for geometric distortions (due to foreshortening) and photometric variations (due to surface mesostructure characteristics.) Second, if we can condense multiple views of the same surface within a single 3-D model, this model could be transmitted very efficiently and rendered under different perspectives and illumination conditions in a virtual environment.

In order to migrate from single-view to multiple-view geology characterization and modeling, we are developing robust and fast algorithms for surface representation and 3-D registration of multiple views, making use of stereo range information. Geometric scene description may also be used to improve texture-based (and possibly shape-based) rock classification. This technique would allow us to improve on current science-based metrics for data representation, ensuring that the hierarchical data segments are assigned transmission priorities and science values using criteria consistent with the objectives of the end users (the scientists) during progressive transmission.

The Telecommunications and Mission Operation Directorate Technology Program is developing new technology for 3-D surface acquisition and for the progressive transmission of shape information. The 3-D geometry of an object can be inferred in different ways depending on the modality of acquisition. If a single stream of images is acquired, then one can derive the motion of the camera along with the structure of the scene represented by the 3-D location of a set of feature points. In the case of a stereo vision system, pairs of images are acquired and range is computed via correlation. The resulting “range images” (Fig. 2) are merged together to infer the 3-D map of the scene (Fig. 3). In both situations we have to deal with the integration of multiple scene descriptions (sparse sets of feature points or collection of surfaces derived from the stereo acquisition) in order to derive a complete representation of the object. The compact description of a 3-D object is based on the inference of a surface representation preserving detected features.

Our surface representation scheme is mesh-based, allowing us to naturally integrate the inferred geometry into CAD/CAM models for finite element studies or visualization purposes. We are studying methods for mesh generation and refinement that preserve the extracted features of the surface. In particular, we are focusing on multiresolution schemes that yield a compact description of a surface.

BRING HOME A ROCK FROM MARS

By representing the 3-D data with a multiresolution model, we will be able to perform a new kind of progressive data transmission based on continuous refinement of the scale of the geometry description of the object. This is a fairly new research area that has already found some application in the progressive transmission of 3-D models on the World Wide Web [3]. We are studying the extension of current progressive compression techniques (such as the ICER algorithm [Kiely 2000] developed by A. Kiely at JPL) to the case of geometric multiresolution descriptions. The 3-D data can be prioritized according to its scientific interest (e.g., a surface of a rock will have higher priority than a surface of sand), and prioritized transmission can be achieved using our hierarchical buffer manage-

CONTINUED ON NEXT PAGE




Figure 2. A stereo image pair of a rock (left and center) and the reconstructed range (right)

ment infrastructure [Kiely 2000, Manduchi 2000]. Indeed, the 3-D geometric surface description will also benefit the geological classification of images. If the surface geometry is known, it can be used to normalize color/multispectral and texture features, thereby increasing the robustness of the classifier.




Figure 3. The 3-D rock shape reconstructed by integrating several stereo image pairs taken from different viewpoints of the rock in Figure 2

Once transmitted, a ground-based researcher can use the geometric descriptions, together with the albedo information extracted from the registered images of the surface, to reconstruct a virtual image of the surface. In particular, by registering and stitching together enough images of a Mars rock taken from different viewpoints, one may be able to render a 3-D virtual image of the rock, which can be manipulated using a 3-D graphics application in a virtual reality environment.

CONCLUSIONS

It might take a while to send the first man to Mars. Meanwhile, we can be content with dwelling in “virtual” 3-D landscapes populated with models of “real” rocks and surfaces collected by rovers on the Martian surface. The technology discussed in this article (encompassing computer vision, autonomous geology, data compression, and graphic rendering) is a contribution toward bringing the geology of the red planet within closer reach of all users—the scientist as well as the general public.

REFERENCES

1. Kiely, A., R. Manduchi, and M. Klimesh, “Maximizing Science Return: Data Compression and On-board Science Processing,” TMOD Technology and Science Program News, Issue 12, June 2000, pp. 12-15.
2. Gillmore, M.S., R. Castano, T. Mann, R.C. Anderson, E.D. Mjolsness, R. Manduchi, and R.S. Saunders, “Strategies for Autonomous Rovers at Mars,” accepted for publication in the *Journal of Geophysical Research-Planets*, 2000.
3. <http://www.viewpoint.com>
4. Manduchi, R., S. Dolinar, F. Pollara, and A. Matache, “Onboard Science Processing and Buffer Management for Intelligent Deep Space Communications,” IEEE Aerospace 2000, Big Sky, Montana, March 2000. 🐱

SPACECRAFT NAVIGATION USING A GRAYSCALE OPTICAL CORRELATOR

TIEN-HSIN CHAO

INTRODUCTION

Recent technical progress has been made in developing a compact high-speed Grayscale Optical Correlator (GOC) for real-time pattern recognition. This GOC, funded partially by the Telecommunication and Mission Operation Technology (TMOT) program at JPL, is being investigated for spacecraft navigation applications. Current status of hardware development, software simulation, and experimental demonstration of real-time landmark tracking during a lander descent sequence will be described.

Many future JPL/NASA flight missions will require on-board navigation and guidance processing to assure a precise and safe landing. One reason for this is that the most scientifically interesting landing site often presents the most formidable challenge for on-board processing. The compact Grayscale Optical Correlator can be used to identify and lock on to the desired landing site to achieve precision landing.

REAL-TIME PATTERN RECOGNITION

The interest in GOC for pattern recognition applications stems from the fact that optical processing inherently offers massive parallelism, ultrahigh speed, and shift invariance, all with large (e.g., 1024×1024 pixel) frame size—not easily achievable by electronics alone. It is particularly advantageous to use a GOC for target recognition when processing a very large frame input. This is due to the fact that the electronic correlator speed will decrease nonlinearly with the increase in input image size, while that of the optical correlator is independent of the input image size. During a lander descent, a larger frame landing site image would provide higher resolution and more landing precision.

The system schematic of the JPL-developed GOC is shown in Figure 1a. In operation, an input image containing the target of interest (TOI) is first fed into the Spatial Light Modulator (SLM).

A collimated laser beam, emanated from a laser diode, is then used to back-illuminate this SLM. Since an SLM operates like an updatable electronic transparency, the throughput laser beam will carry the input image. This image-carrying laser beam will then pass through a Fourier transform (FT) lens that will transform the input into its spatial Fourier spectra at its back focal plane. A second, reflective-type SLM, loaded with filter template, precomputed with training TOI data, is placed in the FT plane. Thus, the image spectra will reflect off this filter SLM and be multiplied with the filter transmittance. The throughput light will then be *inverse* Fourier-transformed by a second FT lens. According to the correlation theory, the correlation result will be found at the back focal plane of the second FT lens. A photo-

CONTINUED ON NEXT PAGE




Figure 1a. GOC system schematic




Figure 1b. Photograph of the portable GOC

detector array (e.g., a charge-coupled device, or CCD) will be placed at the correlation plane to record the output. If one or more targets are matched with that of the correlation filter, a bright correlation peak, corresponding to each target, will be detected at the correlation plane at the centroid of the target location. Thus, any target present in the input data will be easily detected by the presence of the correlation peak at the output plane.

JPL has recently developed, for the first time, a compact, camcorder-sized GOC capable of robust target recognition (fig. 1b). This portable GOC has been successfully field-tested for several real-time target recognition experiments.

A key to successful target recognition using this GOC is the development of an efficient correlation filter algorithm. To date, JPL has developed a distortion invariant correlation filter customized for scale, orientation, and perspective invariant pattern recognition. This filter algorithm, entitled "Maximum Average Correlation Height" (MACH), is capable of computing the optimal trade-offs among several conflicting performance measures, such as Average Correlation Height (ACH), Average Similarity Measure (ASM), Average Correlation Energy (ACE), and Output Noise Variance (ONV).

RECENT STUDIES

TMOT has sponsored the highly compressed MACH correlation filter development to enable the continuous tracking of landing sites with scale and orientation variation, a need specific to JPL/NASA. The rule-based MACH optimization routine enables fast filter computation, and the grayscale MACH filter profile is particularly suitable for optical implementation. An optically implemented MACH filter could be switched at a high speed (1 millisecond/frame) that is three orders of magnitude greater than available digital technology.

Recently, this GOC has been investigated for the application of spacecraft precision landing navigation. We have developed a lander descent

simulation video sequence consisting of images with continuous shift, scale and rotation variations. This sequence was acquired from a high-resolution aerial photograph taken over a volcano. Landing sites contained in each of 65 descent images are clearly detected by the GOC. A total of five MACH filters are used to cover target variations with 1:3 scale changes and 360-degree rotation.

This successful application of MACH correlation filter algorithms has demonstrated a more than 20:1 reduction of data processing as compared to current state-of-the-art template matching methods. This is a significant result for employing the GOC on NASA space applications.

One of the experimental results of the precision landing navigation simulation is shown in Figure 2, which includes a midcourse landing site image and the corresponding correlation peak output, recorded at the output plane of the GOC.

In order to equip future NASA spacecraft with on-board target recognition capability, we continue to miniaturize the GOC to meet stringent spacecraft power and volume constraints. We plan to shrink the GOC volume from its current camcorder-size to matchbox-size. As a key milestone to this process, a miniature 512 × 512 SLM with a much smaller feature size is being developed. By using a pair of these new SLMs, the volume of the GOC will be reduced by more than 100-fold, to less than one cubic inch. In the meantime, JPL has also begun to work with industrial partners to investigate advanced packaging technologies that would be suitable for packaging the miniature GOC for use in a space environment. 🚀




Figure 2. A recent experimental demonstration of spacecraft precision landing navigation using a GOC: left, a landing site image during midcourse descent; right, a correlation peak showing the location of the landing site as identified by the GOC

AUTOMATING AEROBRAKING

DANIEL LYONS

INTRODUCTION

Many years ago, as an undergraduate at the University of Michigan, I studied how to compute the effects of drag on an orbit. After coming to work at JPL, I was extremely fortunate to have the opportunity to apply what I learned in school to the first deliberate use of atmospheric drag to change the orbit of an interplanetary spacecraft during the Magellan mission to Venus. Because the Magellan aerobraking phase was so successful, my next task was to use aerobraking to significantly reduce the propulsive requirements of the Mars Global Surveyor spacecraft in order to launch on a smaller, significantly cheaper launch vehicle. The aerobraking phase of the Mars Global Surveyor mission was also successful, in spite of a broken solar panel. My participation in all of the interplanetary missions which have successfully used aerobraking has taught me that aerobraking is much more challenging and exciting than is taught in school. One particular challenge is automating the aerobraking process to

achieve cost reductions. This article will discuss the technology needed to make aerobraking more robust by using autonomy to eliminate the need for continuous DSN coverage during the aerobraking phase.

AEROBRAKING: What Is It?

Using atmospheric drag to modify an orbit can be divided into two main categories: aerocapture and aerobraking. Aerocapture uses a single pass through the atmosphere to capture the spacecraft and usually requires special thermal protection, such as a heat shield, to enable the spacecraft to survive. Aerobraking uses rocket propulsion to capture the satellite into a highly elliptical orbit, and then uses multiple passes through the atmosphere to achieve a tight, circular orbit. Aerobraking spreads the heating effects over many consecutive orbits so that a heat shield is not required.

CONTINUED ON NEXT PAGE

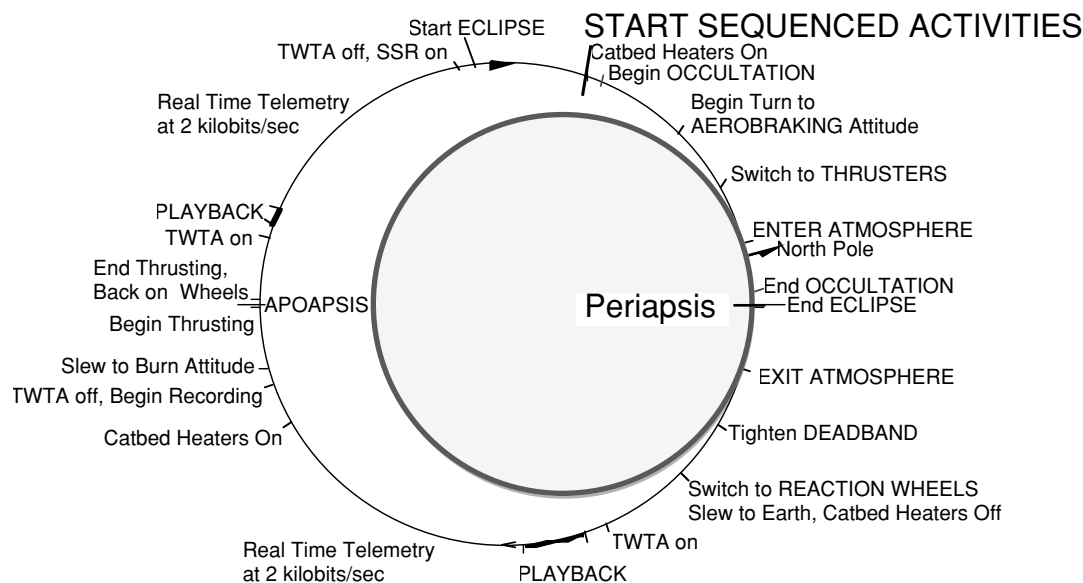


Figure 1: Sequenced activities for a typical MGS orbit

TYPICAL AEROBRAKING ORBITAL EVENTS

Many activities have to occur during a typical aerobraking orbit. Figure 1 shows the sequenced activities during a typical aerobraking orbit from the Mars Global Surveyor mission. Because the aerodynamic torques during the drag pass can be much larger than the control authority of either the reaction wheels or the thrusters, it is very important that the spacecraft be properly configured for the drag pass and commanded to an attitude that will minimize the aerodynamic torques *before* the spacecraft enters the atmosphere. Since the drag pass is nearly centered about periapsis, it is convenient to place activities in the orbit relative to the predicted time of periapsis, even though some events occur before periapsis. (Periapsis is the point on the orbit closest to the center of the planet, and is very close to the point where the atmospheric density and aerodynamic heating rate is highest.)

The activities on-board the spacecraft are specified by a "sequence" of time-tagged commands that are built on the ground and uplinked to the spacecraft. Using the typical MGS orbit in Figure 1 as an example, the first command in the sequence for each aerobraking orbit is the command to turn on the catalyst bed heaters, which warm up the thrusters to minimize damage when they are fired later. The next command is a maneuver using reaction wheels to slew the spacecraft into an attitude close to the desired attitude at atmospheric entry. Any required reconfiguration of the spacecraft occurs at about the same time as the slew maneuver. The spacecraft attitude control mode is switched from reaction wheel control to a thruster control mode with a wide deadband before atmospheric entry so that the attitude control system will not waste propellant trying to fight the aerodynamic torques, which dominate during the drag pass. After exiting the atmosphere, the spacecraft attitude control is switched back to reaction wheel control, the spacecraft is reconfigured, and the attitude is slewed to point the High Gain Antenna at the Earth to play back the data recorded during the drag pass. The beginning and ending times of the playback are specified by the sequence. Some

orbits require a propulsive maneuver near apoapsis to raise or lower periapsis to achieve the desired amount of drag and heating during the next drag pass. The spacecraft remains Earth-pointed until just before the next drag pass, where the sequence of commands for the next orbit repeats the same set of commands. On MGS, the uplinked sequence contained the set of commands for several orbits. Even though the set of commands for each orbit were the same (except for the execution time-tag), the MGS sequence was built as one big list, rather than as a "loop" that called a "subroutine" repeatedly, once for each orbit. (The "looping" approach was used on the Magellan mission, although timing updates were uplinked from the ground, rather than computed on-board.)

Although all interplanetary spacecraft sequences are built on the ground and uplinked, this process is particularly challenging during aerobraking because unpredictable atmospheric fluctuations result in unpredictable changes to the actual times of periapsis. An uplinked sequence that is built on the ground using predicted times rapidly gets "out of phase" with the actual location of the spacecraft on the orbit due to the differences between the predicted and actual drag values on the preceding orbits. Mars Global Surveyor, for example, required a new sequence upload every orbit when the orbit period was larger than about a day. Near the end of the planned aerobraking phase, when the orbit period was only 2 hours, three sequence uploads per day were required just to keep the sequence timing within the 5 minute limits. Multiple orbit predictions become possible for the shorter orbit periods because the change in orbit period is smaller. For an orbit period close to 2 hours, the nominal change in the MGS period per orbit was less than 30 seconds. For the 45 hour orbit period at the start of the MGS aerobraking phase, the planned change in orbit period was several hours per orbit. The 30% 1-sigma atmospheric variability meant that the predicted time of periapsis for the larger period orbits had to be updated after each drag pass so that the sequence would command the spacecraft to the drag attitude within 5 minutes of atmospheric entry.

AUTOMATION OPTIONS

The predicted time of periapsis for all aerobraking missions flown so far has been computed from Navigation orbit reconstructions. The requirement to update the predicted time of periapsis using Navigation orbit reconstructions for MGS required two-way tracking data before and after the drag pass to determine how much the orbit period actually changed. The uncertainty in the time of periapsis, coupled with the need to schedule Deep Space Network (DSN) tracking passes months in advance meant that the DSN antennas had to be allocated to provide continuous coverage during the aerobraking phase. The need to uplink sequences multiple times per day at times that could not be predicted months in advance reinforced the need for continuous tracking. Unfortunately for MGS, the broken solar panel caused the duration of the aerobraking phase to be increased from the planned 4 months to more than 9 months, not counting an additional 7 month hiatus in the middle of the aerobraking phase that was needed for the orbit to achieve the desired orientation to the Sun.

The need for continuous coverage and frequent, critical sequence uploads put a strain not only on the DSN resources, but also the MGS flight team. A more automated sequencing process would have been invaluable during this lengthy aerobraking phase. One of the approaches that is currently being studied to automate the aerobraking process in order to reduce the required two-way tracking coverage and reduce the frequency of sequence uploads is to enable the spacecraft to predict the time of the next periapsis on-board.

ON-BOARD TIMING PREDICTIONS

If the time of periapsis is predicted on-board the spacecraft, a new sequence upload is not required to update just the timing information. In order for this approach to work, the sequences that are uploaded must all be timed *relative* to a common start time that can be linked to the time of the same level of flexibility and control of space-

CONTINUED ON NEXT PAGE




Figure 2. Periapsis Prediction Accuracy using the MGS Accelerometer Data

craft activities as the current process. Only the time-critical task of determining when to start the sequence would be automated on-board the spacecraft. Although MGS had the ability to accommodate relative time-tagged sequences, the project had no confidence in the ability of the spacecraft to compute the periapsis times correctly, no budget to rewrite the flight software to perform the calculation, and no budget to purchase new hardware to eliminate a single-point failure (the accelerometer) in the hardware required to supply the data for the computation. Since the existing flight hardware and software that were inherited from the Mars Observer mission were adequate (as long as continuous DSN tracking was available), the project made the decision to stress the system for a short while during the aerobraking phase.

On-board timing prediction can be divided into two classes: one class that works for any orbit period, and one that only works for shorter period orbits. An approach that will work for any orbit period requires an accelerometer to measure the integral of the deceleration to get the total change in velocity near periapsis. The change in orbit period can be computed by multiplying the total change in velocity times a polynomial function which depends on the orbit period. This computation has been done for all 630 MGS orbits in the MGS accelerometer data archives and the results plotted in Figure 2. Although the MGS accelerometer was not properly scaled for the decelerations experienced during aerobraking (the measurement quantization was determined by another use, and the hardware was inherited), the time of the next periapsis could always be predicted to within 90 seconds using the accelerometer data (on the ground). This 90 second worst-case prediction is well within the 300 second MGS requirement.

Note that the scatter in the prediction accuracy in Figure 2 is much smaller for later orbits where the orbit period is smaller, because the change in the orbit period for a given change in velocity at periapsis becomes much smaller for more circular orbits. Even though the atmospheric variability

remained about the same, the absolute timing prediction error is much smaller for the shorter orbit periods. When the orbit shrinks to the point where the expected period change is equal to the 300 second timing requirement, the time of the next drag pass can be predicted by simply detecting the time of the previous drag pass and then adding the current period (minus 300 sec). For the originally planned MGS mission (before the solar panel failure), the period change per orbit dropped below 300 seconds at an orbit period of about 6 hours. Although the time of the drag pass can be detected using an accelerometer, it can also be inferred from temperature data from surface sensors or from attitude perturbations due to aerodynamic torques. Although this relaxed timing approach would not work well at the start of a typical aerobraking phase when the change in orbit period per orbit is large, it would have been sufficient for the Magellan mission, which started aerobraking from a much lower initial orbit period. (The atmospheric variability at Venus for the Magellan mission was much less than for MGS at Mars, so the change in orbit period was also much more predictable.) Since the timing error must build up over several orbits before it exceeds the timing requirement for the shorter orbit periods, this approach keeps the timing within bounds by making a small timing adjustment every orbit when the orbit period is small.

FULL AUTOMATION

In order to fully automate aerobraking, the spacecraft must be able to autonomously determine when to perform maneuvers to raise or lower the periapsis altitude. Studies are underway to determine how effectively these propulsive maneuvers can be autonomously computed. For aerobraking missions at Mars, it may be possible to automate the maneuvers which control the slow drift in the periapsis altitude, but it will be much more difficult to accommodate the very large changes in atmospheric density that accompany large dust storms on Mars. Dust storm effects are usually inferred from data gathered by scientific


instruments either on-board or on other spacecraft. Since there is not very much data available to characterize dust storm interactions with the atmosphere, developing autonomous algorithms that can predict the effects of a dust storm on the atmosphere will be very challenging, even if instruments are available on-board to supply data in real-time.

SUMMARY

Aerobraking has proven to be a very reliable and robust technique for removing energy from spacecraft orbiting planets with an atmosphere. One drawback associated with aerobraking has been the heavy dependence on continuous DSN coverage. Future aerobraking missions are experiencing greater competition for DSN tracking time as the number of operational spacecraft increases. This report described work that is currently underway to automate as much of the aerobraking process as possible in order to reduce the amount of DSN coverage that is required to support an aerobraking mission. Analysis of the MGS accelerometer data proved that accelerometers could be used to autonomously determine the start time of the sequence of events required for the next aerobraking orbit on-board the spacecraft. Autonomously triggering the sequences would remove the need for continuous, two-way DSN tracking during the aerobraking phase. Autonomously triggered sequences would enable the flight team to spend more time evaluating data to determine spacecraft health because new sequences would only be needed at major events—not multiple times per day!

Aerobraking has been reliably used to change the orbits of both Magellan and Mars Global Surveyor by 1200 m/s. Data from these missions has provided the knowledge needed to automate many of the labor-intensive activities associated with aerobraking. If this knowledge is applied to future missions, it will become possible to aerobrake one spacecraft while simultaneously supporting several other spacecraft at the same planet.

FURTHER READING

- [1] D.T. Lyons, "Aerobraking at Venus and Mars: A Comparison of the Magellan and Mars Global Surveyor Aerobraking Phases", AAS/AIAA Astronautics Conference, Girdwood, Alaska, August 1999. AAS 99-358
- [2] D.T. Lyons, J. Beerer, P. Esposito, M.D. Johnston, W. Willcockson, "Mars Global Surveyor: Aerobraking Mission Overview", *Journal of Spacecraft and Rockets*, Volume 36, Number 3, May-June 1999. pp. 307-313.
- [3] D.T. Lyons, "Mars Global Surveyor: Aerobraking with a Broken Wing", AIAA/AAS Astrodynamics Specialists Conference, Sun Valley, Idaho, August 4-7, 1997, AAS-97-618.
- [4] D.T. Lyons, "Aerobraking: The Key to the Future of Affordable Mars Exploration", Second IAA International Conference on Low-Cost Planetary Missions, John Hopkins University, Applied Physics Laboratory, April 16-19, 1996, IAA-L-0512
- [5] D.T. Lyons, "Aerobraking Magellan: Plan versus Reality", AAS/AIAA Spaceflight Mechanics Meeting, Cocoa Beach Florida, February 14-16, 1994. paper AAS 94-118.
- [6] D.T. Lyons, R. Stephen Saunders, Douglas G. Griffith, "The Magellan Venus Mapping Mission: Aerobraking Operations", 44th Congress of the International Astronautical Federation, Graz, Austria, October 16-22, 1993. paper IAF-93-Q.4.409.
- [7] *Acta Astronautica*, Vol. 35 No. 9-11, pp. 669-676, 1995, Elsevier Science Ltd. Pergamon. 




Figure 1. DSN Block IIA X-Band TWM, top plate and cylinder of the vacuum housing removed. Thermal radiation shields and superconducting magnet have been partially cut away to reveal the maser microwave amplifier module

include a portion of the feed system inside the cryogenic refrigerator with the maser amplifier module. Tops as low as 14 and 15 K have been achieved at X-Band and as low as 24.5 K at Ka-Band. The noise performance of many past DSN maser LNAs and all masers currently in use is shown in Figure 2, which also indicates the year they were introduced.

DSN TWMs provide bandwidth and tuning ranges that far exceed the deep space-to-Earth frequency allocations at S- and X-Band. Wide bandwidth requirements for applications such as continuum radio astronomy are best met with the more recently developed HEMT LNAs, while the cavity maser's very linear characteristics, bandwidth that can be matched to the allocation bandwidth, and the lowest possible noise performance, make these new masers important for the future.

NEW MASER DEVELOPMENTS

The best sensitivity for deep space missions, the highest immunity for radio frequency interference (RFI), and the most robust characteristics that defy burnout, can now be achieved with design techniques recently developed for cavity masers. Additionally, this new technology will require

over an order-of-magnitude less microwave pump power than TWMs, and will contain a much smaller volume of ruby material, allowing use of smaller magnets. The lower pump power requirement and smaller volume will enable these masers to be continuously cooled below 4 K in a CCR, resulting in even lower noise temperatures than achieved by today's best TWMs. Cavity maser module noise temperatures (this is the noise temperature of the amplifier module mounted inside the refrigerator) of 0.6 K at S-Band, 1 K at X-Band, and 2.5 K at Ka-Band can be achieved at a physical temperature of 1.6 K. Slightly higher maser module noise temperatures of 1 K at S-Band, 1.5 K at X-Band, and 3.2 K at Ka-Band will be achieved at physical temperatures of 2.5 K. This noise performance, together with the simplicity of cavity maser construction make cavity masers the logical choice over TWMs for future DSN applications.

The new 4 K Gifford-McMahon (G-M) cryogenic refrigerators (developed by industry) contribute to all current improvements in maser and HEMT performance. These powerful refrigerators provide reliable long-term cooling at a much lower temperature than the physical temperature of about 15 K that was achieved by the previous G-M refrigerators used in most DSN LNAs. Lower physical temperature lowers the maser or HEMT module noise temperature, and also lowers the noise contribution from passive cooled components in front of the module, such as isolators, filters, tracking couplers, and polarizers. G-M refrigerators are being used in all new HEMT LNA designs, and the addition of a helium liquefaction and evaporation stage to these refrigerators will provide the 1.6 to 2.5 K physical temperature preferred by new cavity masers.

HEMTs

HEMT technology was first introduced in 1980; it soon exceeded the microwave amplifier noise performance of the best gallium arsenide (GaAs) field effect transistors (FETs) at both room and cryogenic temperatures. The HEMT is essentially a high-performance field effect transistor whose active region is a finely tuned, multi-layered structure designed to provide a channel for the electron carriers that is almost free of impurities. Thompson-CSF and Fujitsu Laboratories were the first to report the cryogenic noise performance of amplifier modules based on these devices in 1983. In the wake of this success, several other major companies (e.g., GE, Hughes




Figure 2. Noise temperature of DSN masers, 1960–present

Research Labs, TRW, Mitsubishi, NEC) began substantial development programs. Since then, state-of-the-art noise temperatures of cryogenic HEMT-based LNAs have steadily improved. Some significant module noise temperature milestones reached by HEMT amplifiers at various centers worldwide are 5.6 K at 8.5 GHz in 1986, 15 K at 43 GHz in 1993, 30 K at 102 GHz in 1999, 12K at 32 GHz in 1997, 4.5K at 8.4 GHz in 1999, 3.5K at 8.4 GHz in 2000, and 6.5 K at 34 GHz in 2000 (the last four results obtained at JPL).

Beginning in 1987, cryogenically cooled GaAs FET LNAs were installed in the DSN in positions where the ultimate noise performance of the maser was not required, and other features such as lower implementation cost, lower maintenance cost, or wider bandwidths made the FET LNA preferable. With the advent of the HEMT transistor, all existing GaAs FET LNAs were upgraded with HEMT amplifier modules, and more HEMT-based LNAs were introduced into the DSN, totalling over 20 LNAs between 1988 and 1999. The noise performance of past and present DSN FET and HEMT LNAs is shown in Figure 3, which also indicates when they were introduced.

NEW HEMT DEVELOPMENTS

The recent development of indium phosphide (InP) HEMTs is bringing a new level of noise performance to cryogenically cooled, HEMT-based LNAs. A collaboration between JPL and TRW was formed to develop cryogenically coolable InP HEMT devices and MMICs that would outperform other HEMT devices seen to date. The fruits of this collaboration were seen in 1999 with the development of a new InP HEMT device, which resulted in HEMT amplifier modules measuring 4.5 K at 8.4 GHz and 10 K at 32 GHz.

Devices developed in early 2000 by TRW have resulted in the lowest noise HEMT amplifier modules to date: 3.5 K at 8.4 GHz and 6.5 K at 34 GHz. Eight of these X-Band modules have been installed into four dual-channel HEMT LNAs for the DSN's 70-meter antennas. LNA noise temperatures for these four packages measure between 4.0 and 5.5 K. This package is shown in Figure 4. These are the first LNAs implemented in the DSN to use the new 4 K G-M two-stage refrigerators discussed earlier.

This system illustrates a trend in the DSN to move some feed components, especially the

CONTINUED ON NEXT PAGE

Figure 3. Noise temperature of DSN FET and HEMT LNAs, 1987–present.





Figure 4. X-Band dual-channel HEMT package for 70-m antennas. The top plate and cylinder of the vacuum housing and the 30-K thermal radiation shield cylinder have been removed to reveal the cryogenic RF components

polarizer, from the room temperature environment into the LNA's cryogenic package to improve operating noise temperatures. Figure 5 shows the block diagram and system noise budget for the 70-meter antenna at Goldstone, including the previous masers (with all room temperature receive-only feed) and the new dual-channel HEMT (with transmit/receive feed, with some feed functions in the cryogenic package).

A major goal for future LNA implementations in the DSN is to combine: (1) X-Band reception, (2) Ka-Band reception, (3) simultaneous diplexing with a 20-kW X-Band transmitter, and (4) a cryogenic tracking coupler at Ka-Band to aid antenna pointing, into one feed/LNA assembly (whether HEMT or cavity maser). An R&D feed/HEMT LNA package that incorporated all these features was recently tested in one of the DSN's 34-meter antennas. T_e measured 10.5 K at X-Band, 16 K at Ka-Band, and 24 K for the Ka-Band error channel provided by the tracking coupler (in most Ka-Band LNAs and a few X-Band LNAs, the feed system is integrated with the LNA such that LNA noise temperature is not easily measurable. In these cases, T_e is given, which is the input noise

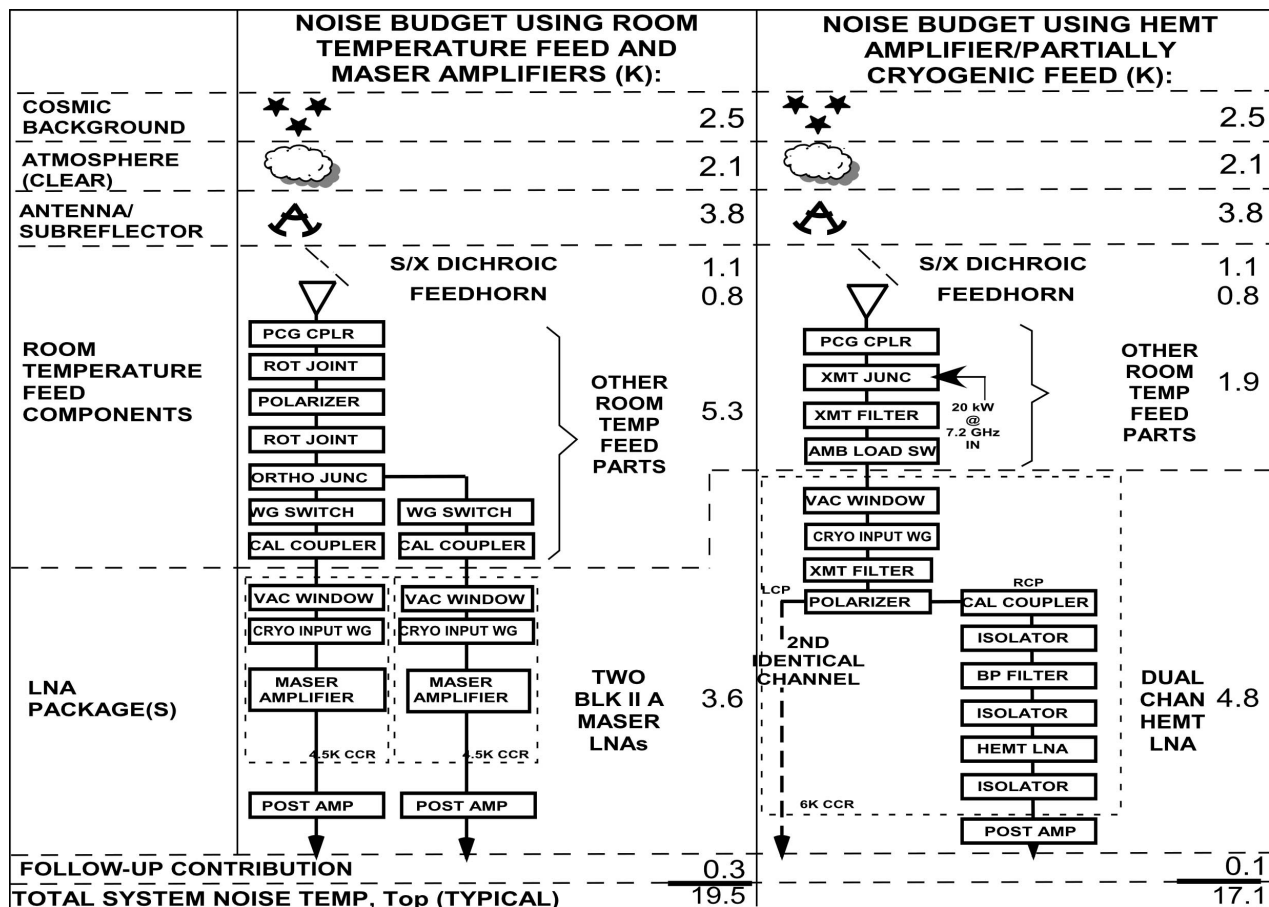


Figure 5. X-Band noise performance of the DSN 70-m antenna at Goldstone with masers with room temperature feed and the new dual-channel HEMT with partially cryogenic feed

temperature measured at the aperture of the feedhorn. T_e is typically 1–3 K higher than LNA noise temperature for these LNAs). This test demonstrated that the goals for this system can be met, and implementation into five DSN antennas is planned over the next four years.

Meanwhile, the collaborative effort is continuing between TRW and JPL. The goal of this work is to thoroughly investigate the device parameters (or trade space) that will yield the best InP HEMT discrete devices and MMICs for cryogenic microwave and millimeter wave low-noise use. The best devices that are currently being tested under this effort should result in a Ka-Band HEMT amplifier module noise temperature of 5 K.

How Much Is a Kelvin Worth?

A 20% (3 K) reduction in the S-Band T_{op} enabled the reception of 117 kbps from Mariner-10 during the Mercury flyby in 1974. This resulted in continuous picture coverage of Mercury along the spacecraft's flight path. Without the T_{op} reduction the mission would have been forced to use a 22 kbps data rate and 80% of the data would

have been lost. An identical system is in use today to support Galileo. It seems that the value of a kelvin grows during critical encounters and when problems occur on spacecraft. X-Band T_{op} reductions of about 20% in 1980 contributed to higher data rates from the Voyager at Saturn, Uranus, and Neptune.

Mission planners have the option of increasing the data rate from a deep-space mission and thereby reducing the tracking time needed to return the same volume of data, or increasing the total data volume received. The percentage reduction in tracking time can be the same percentage by which the T_{op} is reduced. A reduction in the tracking time results in reduced costs to a mission. Use of the time saved to increase the data return enables a greater percentage of data to be returned than the percentage reduction in tracking time (a 20% reduction in T_{op} enables a 25% increase in data rate).

The value of a T_{op} reduction can be based on the costs of using DSN antennas to track deep

CONTINUED ON NEXT PAGE

space missions. Hourly rates for the use of DSN antennas have been estimated for mission planning purposes. These rates, combined for the DSN's three 70-meter and nine 34-meter antennas, were used to show that a 20% reduction in T_{op} (4 kelvin at X-Band) could save about \$20 million per year, or about \$400 million over a 20-year period. Note that a higher antenna utilization time (80% for 7000 tracking hours per year) shows a value of \$280 M for the 20% reduction in T_{op} .

Another way to estimate the value of a T_{op} reduction considers the initial cost and the operations cost of a fully equipped DSN 34-meter beam waveguide antenna over a 20-year period. The construction and implementation cost of approximately \$33 M combined with a \$2 M per year operations cost over a 20-year lifetime gives a total lifetime cost of about \$73 M. Adding five such 34-meter antennas to the DSN increases the DSN's total data reception capability by about 25%. Reducing the X-Band system temperature by 4 K (20%) increases the DSN's total data reception capability by about 25%. It would appear that a 1-K T_{op} reduction at X-Band throughout the DSN is worth about \$91 M and, in this case, a 20% reduction in T_{op} appears to be worth about \$365 M.

The value of a kelvin is frequency-dependent because T_{op} and antenna efficiency are frequency dependent. The DSN's T_{op} at Ka-Band will be higher than at X-Band, perhaps by a factor of two or three, so the value of a 2–3 K improvement will be about \$91 M at Ka-Band.

CONCLUSION

The designers and builders of the early DSN recognized the value of having a low T_{op} . The use of masers was cost-effective because the cost of the tracking station with a large antenna was much more than the cost of the maser and cryogenic cooling system. As more antennas were added to the DSN, and as the DSN moved up in frequency from L-Band (960 MHz) to S-Band, X-Band, and lately Ka-Band, maser amplifiers eventually numbering over 24 systems operated continuously in the DSN.

In the late 1980s, the GaAs FET, followed by the HEMT, became available as a wider bandwidth, lower cost, and reliable LNA option for DSN applications where maser noise performance was not required. InP HEMTs have further narrowed the gap in noise performance between

transistor amplifiers and masers during the last few years, and are now being used in all new X-Band and Ka-Band HEMT receiving systems. Today, 12 masers and 23 HEMTs are operating on 12 DSN antennas.

The importance of a very low T_{op} for tracking deep-space missions has been shown. A new cavity maser design and associated cryogenic refrigerator technology will deliver the lowest practical noise temperatures with T_{op} reductions of about 20% expected for the DSN. This maser will also give the best protection against the growing levels of radio frequency interference (RFI) that are being experienced by the DSN. The combination of HEMTs and cavity masers gives the DSN the best-suited LNAs needed for the challenging requirements of the new millennium.

BIBLIOGRAPHY

Maser

- [1] N.A. Renzetti, et. al., "The Deep Space Network—A Radio Communications Instrument for Deep Space Exploration," JPL Publication 82-104, Jet Propulsion Laboratory, Pasadena, CA, July 15, 1983, Chapter 4.
- [2] James S. Shell, et. al., "Ruby Masers for Maximum G/T_{op} ," Proceedings of the IEEE, Vol. 82, No. 5, May 1994, pp. 796–810 (these first two references describe maser development for the DSN through 1993).
- [3] J. Shell and R.B. Quinn, "A Dual-Cavity Ruby Maser for the Ka-Band Link Experiment," TDA Progress Report 42-116, Jet Propulsion Laboratory, Pasadena, CA, February 15, 1994.
- [4] J.S. Shell and R.C. Clauss, "A 32-Gigahertz Coupled-Cavity Maser Design," TDA Progress Report 42-142, Jet Propulsion Laboratory, Pasadena, CA, August 15, 2000.
- [5] T. Kuriyama, R. Hakamada, H. Nakagome, Y. Tokai, M. Sahashi, R. Li, O. Yoshida, K. Matsumoto, T. Hashimoto, "High Efficient Two-Stage GM Refrigerator with Magnetic Material in the Liquid Helium Temperature Region," Advances in Cryogenic Engineering, Vol. 35, Edited by R.W. Fast, Plenum Press, New York, 1990.

HEMT

- [6] Randall K. Kirschman, ed., "Low-Temperature Electronics," IEEE Press, NY, 1986; Part 4, "Microwave Devices," by S. Petty, pp. 358–363 (this paper references cryogenic GaAs and HEMT amplifier development through 1985).

CONTINUED ON BACK PAGE

THE SMALLEST ANTENNA IN THE DEEP SPACE NETWORK

GEORGE RESCH, C. JACOBS, S. KEIHM, G. LANYI, R. LINFIELD, C. NAUDET, A. RILEY,
H. ROSENBERGER, AND A. TANNER



When you visit the Goldstone complex in March 2001 you will find the DSN's newest—and smallest—dish antennas in operation. They will be placed at equidistant points on the south side of the main DSS-25 pad, like sentinels guarding the scientific treasures we expect to harvest during the radio science experiments planned for the Cassini spacecraft. However, instead of guard duty, these new devices have the job of explaining the terrible things that happen to the radio wave as it travels through our atmosphere, both on the way from the DSN antenna to a spacecraft, and on the way back to Earth. These devices are called water vapor radiometers, or WVRs for short.

A photo of one of the new WVRs, taken at DSS-13 where it was installed temporarily for testing, is shown in Figure 1. With a mass only 0.00003 times that of its 34-m “big brother,” its cost (including development) is 0.02 times greater—so please don’t bump into it! The big white disk is a carefully machined parabolic surface that gathers radio energy from a “beam” in the sky about 1 degree wide and focuses that energy into the white box located at the top of the vertical arm that contains the WVR analog electronics. There are motors that move it around the central pedestal in azimuth, and other motors that move the arm in elevation so the WVR can point in pretty much the same direction as the 34-m antenna.

The primary function of the WVR is to measure temperature. Those of you aware of a DSN 34-m antenna’s cost, given the numbers above, no doubt would lead you to think, “My gosh, this has got to be the world’s most expensive thermometer!” We’re in no position to argue the point, but we will remind you that this is a pretty special thermometer. The WVR measures power, which we translate to brightness temperature; then, using something called a retrieval algorithm we estimate excess path delay due to the presence of water vapor in the direction the WVR points. We do have to keep in mind that occasionally a radio source or planet can creep into the field of view and confuse the results. Another special feature of this thermometer is its combination of sensitivity and stability. It can regularly measure temperature changes of 0.01 kelvin. We can believe these changes are real because the radiometer is stable over time periods of many hours.

These characteristics of the WVR are important because as we track a radio source across the sky (spacecraft or quasar) the total amount of atmosphere the radio wave must traverse varies smoothly, approximately as the cosecant of the elevation angle. However, turbulence in the atmosphere is not smooth and, although small, will cause a slight advance (or retardation) in the phase of the radio wave relative to a path through a vacuum. Such variations in phase of the radio waves we receive can sometimes mask the “signal” that we search for in various radio science and radio astronomy experiments. The fluctuations are dominated by variations in

CONTINUED ON NEXT PAGE



Figure 1. The new advanced WVR with the MTP (center) and the J-series WVR (right).

the amount of water vapor along the path—exactly what the WVR was designed to measure.

How it works

Suppose you were to take a very broadband antenna, amplifier, tunable filter, and power meter outside and point the antenna toward the zenith. If you recorded the power, subtracted the noise power generated by the antenna and amplifier, converted to equivalent blackbody brightness temperature, and plotted the result versus the frequency, then you would obtain something that looked like Figure 2. The bottom plot shows what you would see if there were absolutely no water vapor in the atmosphere. At 1 GHz the temperature is about 4.5K, consisting of 2.7 K from the cosmic blackbody background plus a few Kelvin from oxygen molecules. The temperature increases smoothly along the wings of a complex of spectral lines from the oxygen molecule that peak at around 60 GHz. Adding 2 cm of precipitable water vapor to the atmosphere would yield the top curve in Figure 2. This does not sound like much water vapor, but it corresponds to a fairly humid day at Goldstone and is enough to cause the equivalent of 12 cm of extra delay to a radio wave (at zenith).

The first thing that catches your eye at low frequencies in Figure 2 is the “bump” centered at 22.2 GHz due to a spectral line from the water vapor molecule. The size of the bump indicates a




Figure 2. Zenith temperature of the atmosphere from sea level for 0 and 2 cm of PWV.

first-order estimate of the amount of water vapor and, hence, the extra path delay. Doesn't that sound easy? Hardly worth all the expense?

Unfortunately, or fortunately (depending on where you sit), we have to calibrate the water vapor delay with an accuracy corresponding to a small fraction of a millimeter, but this calculation cannot be determined using the first-order theory. In addition to the WVR, we have added a package of instruments that measure surface conditions (temperature, pressure, relative humidity), and a microwave temperature profiler (MTP) that gives us a measure of the physical temperature of the water vapor as a function of altitude. We won't go into all the gory details here but refer the interested reader to Ref. 1, 2, and 3.

Does it work?

The short answer is “we think so,” but we haven't proved it, at least to the level required by the Gravitational Wave Experiment (GWE) on Cassini. Our goal for the Media Calibration System is to deliver a calibration system whose error budget is known and operates as good or better than the requirement. That way, if the GWE actually detects a gravitational wave, the WVR team gets to bask in the reflected glory (maybe a Nobel Prize?). On the other hand, if there is no detection the investigators can set a meaningful limit on the existence of these waves that every modern theory of gravity predicts but nobody has ever detected. A meaningful limit will constrain many theories regarding the generation of gravitational radiation, and this Cassini experiment will take years to surpass.

When you build an instrument that measures something better than anything else in the world you are faced with a conundrum: How do you prove it? The WVR team faced this conundrum a few years ago and decided we would prove it by calibrating a data type that we thought we knew very well and was similar to the data type that would be used during the Cassini GWE. We made the decision that we would use a connected element interferometer (CEI) between DSS-13 and -15 and see how well we could calibrate the fluctuations in interferometric phase (which we believed to be dominated by water vapor) with the new Media Calibration System.

Figure 3 shows the data that indicates why we think we can meet the Cassini requirements and Figure 4 shows why we haven't completed the job.




Figure 3. The residual delay of the CEI and WVR for a long scan on DOY 138, 2000. The data is site-differenced, mapped to zenith, with linear trends removed.

In Figure 3 we have taken the CEI data and subtracted all the geometric effects as best we could. The solid curve is then the residual CEI error. The dotted curve in Figure 3 is the difference in the WVR data between DSS-13 and -15. The obvious correlation between these data types confirms that CEI residuals are dominated by water vapor fluctuations.

In Figure 4, we go one step further and present the data in the domain of Allan Standard Deviation (ASD), which is the domain in which the radio science requirements are stated. The thin solid and dotted curves show the residual CEI signal and station-differenced WVR data. The

CONTINUED ON BACK PAGE




Figure 4. The Allan Standard Deviation plotted as a function of sampling time for the long scan of DOY 138, 2000. The figure shows the CEI residual data, the WVR residual data, the CEI data after WVR corrections, and the requirements for the Cassini GWE.

dashed curve shows what happens when we use the WVR data to correct the CEI residuals, similar to the way we would correct Doppler data in the Cassini GWE. The improvement is dramatic and meets the requirement on time scales between 2000 and 6000 sec. Other data indicates we can meet the requirement down to a bit less than 1000 sec, but we have been unable to demonstrate we can maintain this out to 10,000 sec.



WHAT NEXT?

We are getting ready for another series of comparison experiments. The TMOD Technology Program has funded the installation of a Stabilized Fiber Optic Distribution Assembly (SFODA) between SPC-10 and DSS 13 that we believe will fix the instrumental problems in the CEI which have prevented us from meeting requirements out to 10,000 sec. With additional support from the Technology Program and the Cassini Radio Science Implementation task for data acquisition and analysis, we will attempt at least four additional experiments under conditions that should be very similar to those we expect during the Cassini GWE.

ACKNOWLEDGEMENTS

We would like to thank Larry Teitelbaum for helpful discussions on data analysis and John Eric Clark, Charles Snedeker, Lyle Skjerve, Leroy Tanida, the staff of DSS-13, and the Operations crews at the Goldstone Signal Processing Center for the invaluable assistance provided during these experiments.

REFERENCES

- [1] Resch, G.M., "Water Vapor Radiometry in Geodetic Applications," in *Geodetic Refraction*, Ed. F.K. Brunner, Springer-Verlag, Berlin, pp 53–84, 1984.
- [2] Westwater, E.R., "Ground-Based Microwave Remote Sensing of Meteorological Variables," in *Atmospheric Remote Sensing by Microwave Radiometry*, Ed. M. Janssen, John Wiley & Sons, New York, 1993.
- [3] Teitelbaum, S.J. Keihm, R.P. Linfield, M.J. Mahoney, and G.M. Resch, "A demonstration of precise calibration of tropospheric delay fluctuations with water vapor radiometers," *Geophys. Research Lett.*, vol. 23, pp. 3719–3722, 1996. 
- [7] M. W. Pospieszalski, et al., "Noise parameters and light sensitivity of low-noise high-electron-mobility-transistors," *IEEE Trans. Electron Devices*, vol. ED-33, pp. 218–223, 1986.
- [8] M. W. Pospieszalski, et al., "Millimeter-wave, cryogenically-coolable amplifiers using AlInAs/GaInAs/InP HEMTs," *Proc. 1993 IEEE MTT-S International Symposium*, Atlanta, GA, pp. 515–518, June 1993.
- [9] S. Weinreb, R. Lai, N. Erickson, T. Gaier, and J. Wielgus, "W-band InP Wideband MMIC LNA with 30 K Noise Temperature," *IEEE MTT-S Digest*, pp. 101–104, 1999.
- [10] M.R. Murti, et al., "Temperature Dependent Small-Signal And Noise Parameter Measurements And Modeling On InP HEMTs," *Proc. 2000 IEEE MTT-S International Symposium*, Boston, MA, June 2000.
- [11] J.J. Bautista, et al., "Cryogenic, X-band and Ka-Band InP HEMT Based LNAs for the Deep Space Network," to be presented at the 2001 IEEE Aerospace Conference, Big Sky, Montana, March 10–17, 2001. 

ON THE WEB

The *TMOD Technology and Science Program News* and related TMOD features are located at:

<http://tmot.jpl.nasa.gov/index.html>

Click on "Program Overview Information," then click on "TMOD Program News."

DISTRIBUTION

To have your name added to or deleted from the *TMOD Technology and Science Program News* distribution list, please call 4-9071.

The *TMOD Technology and Science Program News* is a publication of JPL's Telecommunications and Mission Operations Directorate (TMOD). The TMOD Technology Program is managed by James Lesh and the Science Program by Michael J. Klein.

Managing Editor Charles T. Stelzried
Associate Editor Michael Gregory

JPL

Jet Propulsion Laboratory
California Institute of Technology

JPL D-15493, Issue No. 13, 01/01